

METHOD AND APPARATUS FOR SYNCHRONIZING MULTIPLE VERSIONS OF DIGITAL DATA

TECHNICAL FIELD

The present invention relates to synchronizing multiple versions of multimedia objects by maintaining a history graph and descriptive metadata for the transition between points in the history graph. More particularly, the present invention is directed to a method and apparatus for 5 synchronizing multiple versions of multimedia objects by assigning multimedia objects respective unique identifiers, and by creating a history graph having nodes that store unique identifiers and whose vectors describe the relationship between the multimedia objects.

BACKGROUND OF THE INVENTION

With the advent of networked computers, distributed computing, replicated data storage, mobile computing, and especially the use of multiple computing devices by a single user, there is a problem of managing multiple duplicates of the same object, such as a document or an image. For example, a user may take pictures using a digital camera and then transfer some of the pictures to a mobile computing device, such as a notebook PC. Later, some of the pictures may 15 be transferred to a desktop PC, edited and added to a catalog. This process may be repeated with multiple pictures. In the process, the cataloging software may create thumbnail representations of some of the pictures. A few of the pictures may find their way to a Web site. At every stage, an image may be edited, reformatted, or combined with other images. If at some other time the user wants to consolidate the multiple, somewhat different copies of these pictures, the user is 20 confronted with virtually an impossible task. That is because it is quite difficult to remember or track how those images relate to each other, and very difficult to analyze the differences between

similar images, either automatically or by human intervention. If the user wants to maintain duplicate copies of the “image album” both on the home PC and mobile PC, and the user occasionally applies changes to either one, it is next to impossible to propagate those changes from one location to another. These difficulties also arise with respect to other multimedia objects such as audio and video data.

5 While tools exist to automatically synchronize versions of text documents, no such tools are available for other multimedia object comparison and merging. In addition, the nature of multimedia objects such as image data is much more complicated than pure text, making the problem much more complex to solve.

10 There are cameras available that save the filename as picture 0004, the fourth picture taken in the history of this camera, and the next image will be saved as picture 0005, so there is some level of a unique identifier. However, these identifiers are not unique to all the cameras in the world. For every such camera, the fourth picture taken will be saved as picture 0004.

15 Therefore, there is a need for a method and apparatus for synchronizing multiple versions of multimedia objects by assigning multimedia objects unique identifiers, respectively, and creating a history graph having nodes that store unique identifiers and whose vectors describe the relationship between the multimedia objects. In addition, there is also a need for a method and apparatus for synchronizing multiple versions of multimedia objects by adding metadata associated with each unique identifier describing the transformations applied to the 20 corresponding multimedia object.

SUMMARY OF THE INVENTION

The present invention solves the foregoing deficiencies of the prior art by providing a quick, manageable method and system for synchronizing multiple versions of multimedia objects. More particularly, according to an aspect of the present invention each multimedia object may be identified by a unique identifier, rather than according to their location in a database or a file system. The use of unique identifiers allows synchronization of individual objects and reduces overhead in each system. In addition, a history graph is generated and maintained for each object, where the history table includes nodes that store unique identifiers and whose vectors describe the relationship between the multimedia objects. Metadata may be used to describe the transformations of object. Using metadata to describe the transformations of an object and not physical data changes works more efficiently for multimedia data.

Prior systems required an application specific conflict detection mechanism to determine whether two things were done at the same time in which case there would be a conflict. According to an aspect of the present invention, a substantially unified system is provided by embedding the history graphs and metadata in the actual objects or by maintaining the history graphs in another location, there may be a unified way anyone can tell what manipulations have been performed on the object.

According to an aspect of the present invention, when an object is created, duplicated, moved to a new location, or modified in any way including creating an object by combining multimedia objects, the resulting object is assigned a new identifier. In addition, the unique identifier for each object may be maintained when the object is transferred or edited to facilitate tracking the path of the object.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 a block diagram of a general-purpose digital computing environment that may be
5 used to implement various aspects of the present invention;

FIG. 2 illustrates an example of the evolution of a digital image;

FIG. 3 illustrates an example of manipulating images and updating information according
to an aspect of the present invention;

FIG. 4 illustrates a history graph for the manipulations performed in FIG. 3;

FIGs. 5A-5D illustrate stored images and associated history graphs, respectively, in
accordance with an aspect of the invention;

FIG. 6 illustrates a flowchart for processing an image according to an aspect of the
invention; and

FIG. 7 illustrates a flowchart for determining whether a received image relates to a stored
image(s), according to another aspect of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

In order to facilitate the description of the present invention, the invention will be
described with respect to the particular example of synchronizing images. The present invention,
20 however, is not limited to any particular multimedia object or limited by the examples described
herein. Therefore, the description of the embodiment that follows is for purposes of illustration
and not limitation.

FIG. 1 illustrates an example of a suitable computing system environment 100 on which the invention may be implemented. The computing system environment 100 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computing environment 5 100 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 100.

The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers (PCs), server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above-systems or devices, and the like.

The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that may perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are preformed by remote processing devices that are linked through a communications network. In a distributed computing environment, 15 20 program modules may be located in both local and remote computer storage media including memory storage devices.

Figure 1 is a schematic diagram of an exemplary conventional general-purpose digital computing environment that can be used to implement various aspects of the invention. Computer 100 includes a processing unit 110, a system memory 120 and a system bus 130 that couples various system components including the system memory to the processing unit 110.

5 System bus 130 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. System memory 120 includes a read only memory (ROM) 140 and a random access memory (RAM) 150.

A basic input/output system (BIOS) 160 containing the basic routines that help to transfer information between elements within the computer 100, such as during start-up, is stored in ROM 140. Computer 100 also includes a hard disk drive 170 for reading from and writing to a hard disk (not shown), a magnetic disk drive 180 for reading from or writing to a removable magnetic disk 190, and an optical disk drive 191 for reading from or writing to a removable optical disk 192, such as a CD ROM or other optical media. Hard disk drive 170, magnetic disk drive 180, and optical disk drive 191 are respectively connected to the system bus 130 by a hard disk drive interface 192, a magnetic disk drive interface 193, and an optical disk drive interface 194. The drives and their associated computer-readable media provide nonvolatile storage of computer readable instructions, data structures, program modules and other data for personal computer 100. It will be appreciated by those skilled in the art that other types of computer readable media which can store data that is accessible by a computer, such as magnetic cassettes, 20 flash memory cards, digital video disks, Bernoulli cartridges, random access memories (RAMs),

read only memories (ROMs), and the like, may also be used in the exemplary operating environment.

A number of program modules can be stored on the hard disk, magnetic disk 190, optical disk 192, ROM 140 or RAM 150, including an operating system 195, one or more application programs 196, other program modules 197, and program data 198. A user can enter commands and information into computer 100 through input devices, such as a keyboard 101 and a pointing device 102. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 110 through a serial port interface 106 that is coupled to the system bus, but may be connected by other interfaces, such as a parallel port, a game port or a universal serial bus (USB). A monitor 107 or other type of display device is also connected to system bus 130 via an interface, such as a video adapter 108. In addition to the monitor, personal computers typically include other peripheral output devices (not shown), such as speakers and printers.

Computer 100 can operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 109. Remote computer 109 can be a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to computer 100, although only a memory storage device 111 has been illustrated in Figure 1. The logical connections depicted in Figure 1 include a local area network (LAN) 112 and a wide area network (WAN) 113. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, computer 100 is connected to local network 112 through a network interface or adapter 114. When used in a WAN networking environment, personal computer 100 typically includes a modem 115 or other device for establishing a communications over wide area network 113, such as the Internet. Modem 115, 5 which may be internal or external, is connected to system bus 130 via serial port interface 106. In a networked environment, program modules depicted relative to personal computer 100, or portions thereof, may be stored in the remote memory storage device.

It will be appreciated that the network connections shown are exemplary and other techniques for establishing a communications link between the computers can be used. The existence of any of various well-known protocols, such as TCP/IP, Ethernet, FTP, HTTP and the like, is presumed, and the system can be operated in a client-server configuration to permit a user to retrieve web pages from a web-based server. Any of various conventional web browsers can be used to display and manipulate data on web pages.

As noted above, although the description of the invention refers to digital images, it is to be understood that the invention may be used to synchronize versions of other forms of digital data such as emails, slide presentations, schematics, operating system executables, audio, video, etc. The present invention provides a quick, manageable solution to synchronizing versions of objects by maintaining a history graph and descriptive metadata for the transition between points on the history graph.

20 In today's digital world, many images are obtained with digital cameras. Often, users find that they have acquired many images or pictures that are similar. This usually occurs because the original picture is modified and/or edited in various ways. For example, the

proliferation of an image may begin when the user sends the image out to different recipient over the Internet. Many users enjoy sharing pictures with friends and relatives by sending them over the Internet. Recipients may manipulate the pictures by cropping, editing, etc. The image may then be returned to the user. This usually results in the user having an original and one or more variants of the original.

At some point, the user may wish to consolidate the images he has accumulated. However, this often proves to be a difficult task as it is often nearly impossible to determine how one image differs from another. More particularly, it is almost impossible to determine when the image corresponds to a previous image processed for red-eye reduction or processed for lower resolution, or processed with one of the many other manipulations that can be done to pictures by various products such as Picture It!® by Microsoft®. While computer analysis technology exists for comparing textual documents and determining the differences in the documents, no such technology exists for comparing other multimedia objects, such as images, and determining how similar images are different from each other. Ultimately, users usually end up saving all of the pictures because the user cannot determine which one is better or which one to save.

According to an aspect of the present invention, a quick, manageable solution to the synchronization problem is provided by maintaining, for each image, a history graph and descriptive metadata for the transition between points on the history graph. According to the present invention, a unique identifier may be assigned to each digital image. In addition, a history graph may be created and maintained for each image. The history graph includes nodes that store unique image identifiers, and vectors that describe the relationship between images.

Each vector may have associated metadata describing the transformation from one node to the other.

Upon creation of an image (e.g., image capture by digital camera) or when receiving an image that has not been managed (i.e., from a legacy device), a unique identifier is assigned to the image. In addition, a unique identifier may be generated every time a modification is made to the image. The unique identifier may be generated by any known method, including implicit derivation from image data through methods such as hashing or cyclic redundancy checking (CRC). More particularly, when an image is created, duplicated, moved to a new location, or modified in any way including creating an image by combining other images, the resulting image is assigned a unique identifier. In order to facilitate tracking the image path, the unique identifier is not deleted or modified when the image is transferred or edited. The unique identifier may be a Global Unique Identifiers (GUID). GUIDs are usually easy to generate and large enough to support unique identifiers. The unique identifier may be used in combination with the camera serial number and/or manufacturing code like UPC. Depending upon the specifics of the implementation of the invention on a particular device, the unique identifier may either be stored with the image (e.g., file system that supports extended file attributes, image file format that supports association of metadata with the image such as EXIF: Exchangeable Image File Format), or in a separate database. In the later case, a pointer to the location of the image may be stored together with the unique identifier.

According to the present invention, metadata may be associated with each unique identifier. The metadata describes the way in which the new image differs from the image(s) from which it was created. Metadata usually includes time of creation and location of the object,

such as an image. It may also include a description of the transformation used to derive this image from its predecessor. For example, an image that was modified for display on a TV monitor from a Desktop PC may have the following metadata: Location Pictures/For TV/Vacation; Time 7/9/2001 Brightness +7; Size 1280X960 ≥ 640X480. Additional data beyond identifier and date/time stamp depends on the media. For example, if the object is an image, the user may have the resolution of the image as an example of metadata. This may be generated by the application that was used to create the change. Copying a file will usually not be an operation that would be stored in metadata because it does not result in the manipulation of the data. If the user sends an image to a person via email, the recipient may modify the image and put the image on the Web. The user may then find the image on the Web and determine what changes were made to the image. More particularly, the user may compare the metadata for the image to determine whether the image put on Web is the image received.

In addition to metadata, a history graph having nodes corresponding to the unique identifiers may provide a complete description of the evolution of the image. Metadata may be associated with each vector describing the transformation of the image from one node in the history graph to another. Thus, the relationship between two images may be determined by observing their unique identifiers and history graphs. For example, two images may be considered identical when they have the same unique identifiers, while two images may be related when their history graphs intersect. An example of image evolution is shown in FIG. 2.

According to an aspect of the invention, the history graph for an image may be truncated to save storage space. More particularly, portions of the history graph may be stored with each image, or the history graphs may be stored in a central location. More particularly, partial

history graph and metadata, describing recent history, may be transferred with the image. In other instances, only the partial history graph of recent changes may be transferred. The operating system may determine whether to truncate the history graph/metadata. The system may also be arranged to truncate the history graph/metadata based upon user input.

5 The history graph may be stored with the image or in a central location. For example, if the user is designing an image with a program such as Picture It!® by Microsoft®, it may be more efficient to only deal with pictures that have been imported into Picture It!® and modified with Picture It!®. In that case, the user may want a central directory. Central directories are usually faster to search. On the other hand, if the user plans on sending images to photo processing sites and posting pictures to Web sites, the user may want to store information with the picture.

10 In addition, the history graph may be stored both with the image and in a central location. The duplication of information may prevent the requirement of having to search graphs and compare data. In this case, there is a duplicate in each of the user's applications. In another arrangement, a central service may be used. If a central service is used, then an Application Programming Interface (API) that other programs use whenever they retrieve or store digital data may be needed to interface to the central service.

15 Each device that employs the synchronization method according to the present invention, includes a program that manages image storage and synchronization. The program is usually part of the Operating System (OS) of the device, in the form of a system service or integrated into the device's file system or other storage system. The program that implements the method works in coordination with other software that manipulates digital images. The other software

includes copy, transmit, image editing, synchronization and other programs. The program according to the present invention may provide an API to retrieve or store digital images. Because of the uniqueness of the identifiers associated with the images, the history graph and metadata may be stored separately from the images. In addition, multiple related images may 5 share a single history graph.

In another implementation, the program that implements the method according to the present invention may be integrated with the OS services such that other software that manipulates digital images does not have to coordinate with it explicitly. Rather, the OS takes care of creating and associating the unique identifiers with the multiple digital images. In such a case, the OS API that handles data storage and retrieval may be equipped to receive the metadata associated with each new image. Usually, that is achieved by usage of file formats that embed the metadata inside of the image file.

The history graph and metadata of an image may be used for many purposes in addition to version synchronization. For example, an image's history may be examined by an editing tool to determine whether the image has had representations that may not be compatible with its new representation. More particularly, an image may be transferred from a Desktop PC to a mobile computing device such as personal digital assistant (PDA). Since mobile computing devices often have a much lower screen resolution than the Desktop PC, and also less storage space, it makes sense to create a lower resolution version of the image to be stored on such mobile computing devices. Later, the user of the mobile computing device may attempt to edit the image. The editing tool may examine the history of the image and inform the user that a copy of the image exists on their Desktop PC and that the changes applied to the image on the mobile 20

computing device may not be transferable back to the copy residing on the Desktop PC. Alternatively, a merge tool may be able to understand the type of change and apply it to the other copy of the image (e.g., removing blemish at a specified location of the image).

As an example, a user may find a situation where an image is received and the user
5 already has a catalog and wants to know whether this image already exists. This information may be obtained by reviewing the change history in the history graph that is received together with the image and trying to find whether any of the nodes of the graph correspond to information in the user's history graph. If it does, then the user can determine whether to add this image or to throw it out.

More particularly, suppose the user obtains a JPEG image off of the Internet, and it has a unique identifier and some time/date stamp and maybe resolution. Then the user puts the image into a program such as Microsoft's Picture It!®. The first thing that the program may do is convert the image to its own format. For the purposes of this discussion, assume that the moment that the program creates the image internally in its own format, the program creates a new graph for the image and adds to the history of the changes for that image the fact that it was converted from JPEG into a particular format. A unique identifier and metadata may be generated to indicate how the image data was transformed. Other manipulations may be performed on the image. For example, the image may be cropped, the contrast or balance may be adjusted. These changes may be represented as global changes to the image. When the image
20 is saved, all of the information for the format of the image with the metadata and the new grid and probably the previous incarnation of the image may be saved. Therefore, the GUID of the parent image may be stored. Suppose that the user takes the image and converts it back to the

JPEG format. At this point, a new GUID will be generated and the information that the image was converted from a particular format to JPEG will be added. At some point, the user may post the image on the Web. The user who posted the image on the Web originally may later pick this image up and immediately determine that this image was derived from their original image. This 5 is because somewhere in the list of GUIDs is the GUID that the original use had. Then the user can actually determine what happened to the image later by referring to the metadata.

A specific example will be described with reference to FIGs. 3 and 4. In FIG 3, the processing evolution of an image is shown. The image having GUID1 may be processed for redeye to arrive at the image having GUID2. In addition, the image having GUID3 may be cropped to arrive at the image having GUID4. Finally, the images with GUID2 and GUID4 may be combined to form an image with GUID5. The history graph corresponding to this image processing is shown in FIG. 4. The evolution of the image with GUID5 may be determined from the history graph shown in FIG. 4. The history graph shown in FIG. 4 and the metadata may be transferred together with the image having GUID5 so that the recipient may determine the evolution of the image. The history graph and metadata for an image are not visible upon display. However, a program may read the information in the file and use it.

Another example is illustrated in FIGs. 5A-5D. In this example, the history graph may be stored in a central storage 530. The metadata may be stored separately from the history graph. In this case, the identifier coordinates the information. Images are already stored and they are 20 identified by numbers 57, ..., etc., as shown in FIG. 5A. The history graph indicates that image 68 was derived from image 59, as well as illustrating other details. Referring to FIG. 5C, assume that a new image is received and it has the following information: identifier 71 and short

history that indicates that it was derived from image 67 and that image 67 was derived from image 65. In this arrangement, the image arrives and the system services determine where it fits with respect to other stored images. It performs this operation by looking for the identifier numbers in the history graphs stored in the central storage 530. Once it is determined where the 5 image fits with respect to other stored images, the history graph is updated to include the image, as shown in FIG. 5D. All of the other information for the image may be removed and stored in another location. The actual stored image is different from the received image in the sense that the metadata changed based on change in history. The identifier may be used to find the data in other storage locations.

10 FIG. 6 illustrates processing upon receipt of an image according to an aspect of the present invention. In step S1, an image is received. In step S2, it is determined whether the image was just captured or whether it was received from another source. If the answer in step S2 is Yes, then step S3 is performed and a unique identifier is generated for the image. If the answer in step S2 is No, then step S5 is performed to determine whether the image has an associated unique identifier, metadata, and/or history graph. If the answer in step S5 is Yes, the processing proceeds to step S4. If the answer in step S5 is No, then a unique identifier is assigned to the image in step S6. In step S4, is determined whether the received image was modified/combined with stored image(s) after being received. If the answer in step S4 is Yes, then, appropriate metadata describing the transformations or manipulations performed on the 20 received image is added to the data field and a history graph is created for the image in step S8. Finally, the resulting image and the metadata/history graph are stored in step S9. If the answer in step S4 is No, then processing proceeds to step S9 and the image is stored together with its

unique identifier. Once a unique identifier has been assigned in step S6, step S7 is performed to determine whether the received image was modified/combined with stored image(s) after being received. If the answer in step S7 is Yes, then, appropriate metadata describing the transformations or manipulations performed on the received image is added to the data field and 5 the history graph for the image is updated in step S8. The resulting image and the metadata/history graph are then stored in step S9.

10
15
20
25

Referring to FIG. 7, a process for determining whether a received image corresponds to a stored image is illustrated according to an aspect of the present invention. In step S10, an image is received. In step S11, it is determined whether the image includes a unique identifier, metadata and history graph. If the answer in step S11 is Yes, then processing continues to step S13. If the answer in step S11 is No, then a unique identifier is assigned to the image and the image may be stored in step S12. In step S13, the unique identifier, metadata and history graph information are used to determine whether the received image corresponds to or is related to any of the stored images. In step S14, it is determined whether the comparison in step S13 has discovered an image that is related to or the same as a stored image. If the answer in step S14 is Yes, then processing continues to step S16, where it is determined whether or not the user wishes to store the image or discard it. If the answer in step S14 is No, then the image may be stored with the associated metadata and history graph information in step S15.

As demonstrated by the foregoing, the present invention provides a quick, manageable 20 method and system for synchronizing multiple versions of multimedia objects. More particularly, according to an aspect of the present invention each multimedia object may be identified by a unique identifier. The use of unique identifiers allows synchronization of

individual objects and reduces overhead in each system. In addition, a history graph may be generated and maintained for each object, where the history graph includes nodes that store unique identifiers and whose vectors describe the relationship between the multimedia objects. Metadata may be used to describe the transformations of object.

5 While particular embodiments of the invention have been shown and described, it is recognized that various modifications thereof will occur to those skilled in the art without departing from the spirit and scope of the invention. More particularly, it will be appreciated by those of ordinary skill in the art that the present application may be useful in many applications. For example, the invention may be implemented in an image cataloging system or in a system for tracking multimedia properties of objects. The described embodiments are to be considered in all respects only as illustrative and not restrictive. Therefore, the scope of the herein-described invention shall be limited solely by the claims appended hereto.

10